

GAPS/CoGAPS Users Manual

Elana J. Fertig (ejfertig@jhmi.edu) and Michael F. Ochs (ochsm@tcnj.edu)

May 3, 2016

Contents

Chapter 1

Introduction

Gene Association in Pattern Sets (GAPS) infers underlying patterns in a matrix of measurements that can be interpreted as arising from the multiplication of two lower dimensional matrices. The first development of this code in R/Bioconductor was focused on gene expression analysis, however the original concept was used in spectral imaging. The approach is a general form of matrix factorization using a stochastic algorithm. In this vignette, we will focus on gene expression analysis for concreteness, but the factorization is applicable more broadly.

The Markov chain Monte Carlo (MCMC) matrix factorization that infers patterns also infers the extent to which individual genes belong to these patterns. The CoGAPS algorithm extends GAPS to infer the coordinated activity in sets of genes for each of the inferred patterns based upon (?) and to refine gene set membership based upon (?).

The GAPS algorithm is implemented in C++ and compiled and integrated into R using the Rcpp package. GAPS is licensed under the GNU General Public License version 2. You may freely modify and redistribute GAPS under certain conditions that are described in the top level source directory file `COPYING`.

The R package CoGAPS is designed to facilitate the corresponding analysis of microarray measurements by calling the GAPS C++ library. With the migration to C++ code, the installation as noted in Chapter ?? should now be automatic. Running instructions for the GAPS and CoGAPS analyses are provided in Sections ?? and ??, respectively. CoGAPS and GAPS are freely available at <https://github.com/ejfertig/CoGAPS> and through the CoGAPS Bioconductor package.

If you use the CoGAPS package for your analysis please cite: (?) EJ Fertig, J Ding, AV Favorov, G Parmigiani, and MF Ochs (2010) CoGAPS: an R/C++ package to identify patterns and biological process activity in transcriptomic data. *Bioinformatics* **26**: 2792-2793.

To cite the CoGAPS algorithm use: (?) MF Ochs (2003) Bayesian Decomposition in *The Analysis of Gene Expression Data: Methods and Software* G Parmigiani, E Garrett, R Irizarry, and S Zeger, ed. New York: Springer Verlag.

To cite the gene set statistic use: (?) MF Ochs, L Rink, C Tarn, S Mburu, T Taguchi, B Eisenberg, and AK Godwin (2009) Detection of treatment-induced changes in signaling pathways in gastrointestinal stromal tumors using transcriptomic data. *Cancer Research* **69**: 9125-9132.

To cite the set-membership refinement statistic use: (?) EJ Fertig, AV Favorov, and MF Ochs (2012) Identifying context-specific transcription factor targets from prior knowledge and gene expression data. *2012 IEEE International Conference on Bioinformatics and Biomedicine*, B310, *in press*.

Please contact Elana J. Fertig ejfertig@jhmi.edu or Michael F. Ochs ochsm@tcnj.edu for assistance.

Chapter 2

Installation Instructions

The GAPS and CoGAPS algorithms are implemented in an open source C++ software and an R package to facilitate analysis, visualization, and integration with other R tools (CoGAPS, available through Bioconductor).

With this version of CoGAPS, installation should be automatically completed through Bioconductor package installation:

```
source("http://www.bioconductor.org/biocLite.R")
biocLite("CoGAPS")
```

The C++ software will be automatically compiled, linking to required boost libraries distributed as part of the CRAN BH package (<http://cran.r-project.org/web/packages/BH/index.html>) and Rcpp (<http://cran.r-project.org/web/packages/Rcpp/index.html>).

Chapter 3

Executing CoGAPS

In this chapter, we describe how to run both the GAPS and CoGAPS algorithms.

3.1 GAPS

GAPS seeks a pattern matrix (\mathbf{P}) and the corresponding distribution matrix of weights (\mathbf{A}) whose product forms a mock data matrix (\mathbf{M}) that represents the expression data \mathbf{D} within noise limits (ϵ). That is,

$$\mathbf{D} = \mathbf{M} + \epsilon = \mathbf{A}\mathbf{P} + \epsilon. \quad (3.1)$$

The number of rows in \mathbf{P} (columns in \mathbf{A}) defines the number of biological patterns that GAPS will infer from the measured microarray data or equivalently the number of nonorthogonal basis vectors required to span the data space. As in the Bayesian Decomposition algorithm (?), the matrices \mathbf{A} and \mathbf{P} in GAPS are assumed to have the atomic prior described in (?). In the GAPS implementation, α_A and α_P are corresponding parameters for the expected number of atoms which map to each matrix element in \mathbf{A} and \mathbf{P} , respectively. The corresponding matrices \mathbf{A} and \mathbf{P} are found by MCMC sampling.

3.1.1 Methods

The GAPS algorithm is run by calling the `gapsRun` function in the CoGAPS R package as follows:

```
> gapsRun(D, S, ABins = data.frame(), PBins = data.frame(), nFactor = 7,
  simulation_id = "simulation", nEquil = 1000, nSample = 1000,
  nOutR = 1000, output_atomic = "FALSE", fixedBinProbs = "FALSE",
  fixedDomain = "N", sampleSnapshots = "TRUE", numSnapshots = 100,
  alphaA = 0.01, nMaxA = 1e+05, max_gibbmass_paraA = 100,
  alphaP = 0.01, nMaxP = 1e+05, max_gibbmass_paraP = 100)
```

Input Arguments The inputs that must be set each time are only the data and standard deviation matrices, with all other inputs having default values. However, in reality it is essential to set at least `nFactor`, `nEquil`, and `nSample` based on the expected dimensionality of the data and the number of iterations needed to explore the posterior distribution. The arguments are as follows:

D The matrix of `m` genes by `n` samples of expression data. The input should be a matrix object and row names and column names will be retained in the output matrices as appropriate.

S The matrix of `m` genes by `n` samples of standard deviations for the expression data. The values will be used in estimating the Likelihood, presently under an assumption of normal error distribution.

ABins a matrix of same size as `A` which gives relative probability of that element being non-zero (optional)

PBins a matrix of same size as **P** which gives relative probability of that element being non-zero (optional)

nFactor number of patterns (basis vectors, metagenes), which must be greater than or equal to the number of rows of **FP**

simulation_id name to attach to atoms files if created

nEquil number of iterations for burn-in

nSample number of iterations for sampling

nOutR how often to print status into R by iterations

output_atomic whether to write atom files (large)

fixedBinProbs Boolean for using relative probabilities given in **Abins** and **Pbins**

fixedDomain character to indicate whether **A** or **P** is domain for relative probabilities

sampleSnapshots Boolean to indicate whether to capture individual samples from Markov chain during sampling

numSnapshots the number of individual samples to capture

alphaA sparsity parameter for **A** domain

nMaxA PRESENTLY UNUSED, future = limit number of atoms

max_gibbmass_paraA limit truncated normal to max size

alphaP sparsity parameter for **P** domain

nMaxP PRESENTLY UNUSED, future = limit number of atoms

max_gibbmass_paraP limit truncated normal to max size

The algorithm returns in the results a list with

Amean A matrix with the sampled mean value for the amplitude matrix **A**.

Asd A matrix with the sampled standard deviation for the amplitude matrix **A**.

Pmean A matrix with the sampled mean value for the pattern matrix **P**.

Psd A matrix with the sampled standard deviation for the pattern matrix **P**.

atomsAEquil A vector with the number of atoms in the **A** domain throughout the equilibration steps.

atomsASamp A vector with the number of atoms in the **A** domain throughout the sampling steps.

atomsPEquil A vector with the number of atoms in the **P** domain throughout the equilibration steps.

atomsPSamp A vector with the number of atoms in the **P** domain throughout the sampling steps.

chiSqValues A vector with the sample chi-squared values throughout sampling.

meanChi2 χ^2 value of the final mean result, i.e. $\chi^2 = \text{frac}(D - AP)^2 \sigma^2$.

ASnapshots Samples of **A** matrices taken during sampling.

PSnapshots Samples of **P** matrices taken during sampling.

Once the GAPS algorithm has been run, the inferred patterns and corresponding amplitude can be displayed using the `plotGAPS` function as follows:

```
> plotGAPS(Amean, Pmean, outputPDF="")
```

where setting `outputPDF` to a string will redirect output to a PDF file from the screen. The command

```
> plotP(Pmean, Psd)
```

will create a plot of just the rows of the \mathbf{P} matrix with standard errors, which is equivalent to plotting the basis vectors for the matrix factorization.

Input Arguments

Amean The amplitude matrix **Amean** obtained from GAPS.

Pmean The pattern matrix **Pmean** obtained from GAPS.

outputPDF Name of an pdf file to which the results will be output. (Optional; default="" will output plots to the screen.)

Psd The standard deviation of the mean for **Pmean**.

Side Effects

- Save the plots of **Amean** and **Pmean** to the pdf file `outputPDF`.

3.1.2 Example: Simple Simulation

In this example, we have simulated data in `SimpSim` (`SimpSim.D`) with three known patterns (`SimpSim.P`) and corresponding amplitude (`SimpSim.A`) with specified activity in two gene sets (GSets). In this data set, each gene set is overexpressed in one the simulated patterns and underexpressed in one.

```
> library('CoGAPS')
> data('SimpSim')
> nIter <- 5000
> results <- gapsRun(SimpSim.D, SimpSim.S, nFactor=3,
+                   nEquil=nIter, nSample=nIter)
```

```
Equil:1000 of 5000, Atoms:58(64) Chi2 = 1964.61
Equil:2000 of 5000, Atoms:66(71) Chi2 = 1494.34
Equil:3000 of 5000, Atoms:59(77) Chi2 = 1484.55
Equil:4000 of 5000, Atoms:70(86) Chi2 = 1458.63
Equil:5000 of 5000, Atoms:70(85) Chi2 = 1471.82
Samp: 1000 of 5000, Atoms:66(90) Chi2 = 1481.84
Samp: 2000 of 5000, Atoms:62(87) Chi2 = 1484.46
Samp: 3000 of 5000, Atoms:63(91) Chi2 = 1484.04
Samp: 4000 of 5000, Atoms:64(88) Chi2 = 1477.49
Samp: 5000 of 5000, Atoms:64(97) Chi2 = 1451.79
*** Check value of final chi2: 1451.79 ****
```

```
> plotGAPS(results$Amean, results$Pmean, 'ModSimFigs')
```

```
null device
1
```

Figure ?? shows the results from plotting the GAPS estimates of **A** and **P** using `plotGAPS`, which has a fit to **D** of $\chi^2 = 1383.62346252315$. Figure ?? displays the true patterns used to create the `SimpSim` data, stored in `SimpSim.P`.

Figure ?? shows the results from running `CoGAPS` on the `SimpSim` data.

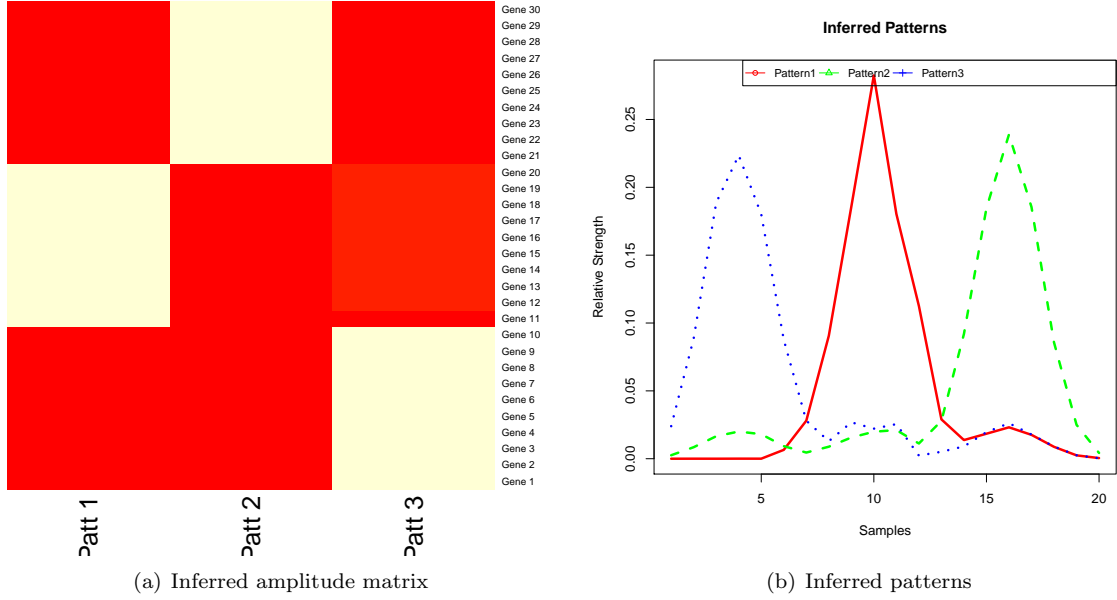


Figure 3.1: Results from GAPS on simulated data set with known true patterns.

3.2 CoGAPS

CoGAPS infers coordinated activity in gene sets active in each row of the pattern matrix \mathbf{P} found by GAPS in a single step, by running both GAPS and then performing the statistical analysis of `calcCoGAPStat`. Specifically, CoGAPS computes a Z -score based statistic on each column of the \mathbf{A} matrix developed in (?). The resulting Z -score for pattern p and gene set i , \mathcal{G}_i , with G elements is given by

$$Z_{i,p} = \frac{1}{G} \sum_{g \in \mathcal{G}_i} \frac{\mathbf{A}_{gp}}{\mathbf{Asd}_{gp}} \quad (3.2)$$

where g indexes the genes in the set and \mathbf{Asd}_{gp} is the standard deviation of \mathbf{A}_{gp} obtained from the MCMC sampling in GAPS. CoGAPS then uses random sample tests to convert the Z -scores from eq. (??) to p values for each gene set.

3.2.1 Methods

The CoGAPS algorithm is run by calling the `CoGAPS` function in the CoGAPS R package as follows:

```
> results <- CoGAPS(data, unc, ABins = data.frame(), PBins = data.frame(),
  GStoGenes, nFactor = 7, simulation_id = "simulation",
  nEquil = 1000, nSample = 1000, nOutR = 1000,
  output_atomic = FALSE, fixedBinProbs = FALSE,
  fixedDomain = "N", sampleSnapshots = TRUE,
  numSnapshots = 100, plot = TRUE, nPerm = 500,
  alphaA = 0.01, nMaxA = 1e+05,
  max_gibbmass_paraA = 100, alphaP = 0.01,
  nMaxP = 1e+05, max_gibbmass_paraP = 100)
```

Input Arguments




Figure 3.2: Known true patterns used to generate SimpSim data.

... Input arguments from GAPS.

GStoGenes List or data frame containing the genes in each gene set. If a list, gene set names are the list names and corresponding elements are the names of genes contained in each set. If a data frame, gene set names are in the first column and corresponding gene names are listed in rows beneath each gene set name.

nPerm Number of permutations used for the null distribution in the gene set statistic. (optional; default=500).

plot Use `plotGAPS` to plot results from the run of GAPS within CoGAPS.

List Items in Function Output

... Output list from GAPS except vectors of sample atom numbers and individual χ^2 values.

D The input data matrix.

Sigma The input standard deviation matrix.

GSUpreg p-values for upregulation of each gene set in each pattern.

GSDownreg p-values for downregulation of each gene set in each pattern.

GSActEst Conversion of p-values to activity estimates of each gene set in each pattern (see (?) for details on conversion. Essentially values near 1 indicate high activity for the set and near -1 indicate low activity.

The CoGAPS algorithm can also be run manually by first running the GAPS algorithm described in Section ?? and then calling the function `calcCoGAPSStat` as follows:

```
> calcCoGAPSStat(Amean, Asd, GStoGenes, numPerm=500)
```

The input arguments for `calcCoGAPSStat` are as described in the previous sections. This function will output a list containing `GSUpreg`, `GSDownreg`, and `GSActEst`.

3.2.2 Example: Simulated data

In this example, we have simulated data in `SimpSim` with three known patterns (`SimpSim.P`) and corresponding amplitude (`SimpSim.A`) with specified activity in two gene sets (`GSets`). In this data set, each gene set is overexpressed in two of the simulated patterns and underexpressed in one.

```
> library('CoGAPS')
> data('SimpSim')
> nIter <- 5000
> results <- CoGAPS(data=SimpSim.D, unc=SimpSim.S,
+                   GStoGenes=GSets,
+                   nFactor=3,
+                   nEquil=nIter, nSample=nIter,
+                   plot=FALSE)
```

```
Equil:1000 of 5000, Atoms:60(55)  Chi2 = 2511.45
Equil:2000 of 5000, Atoms:64(70)  Chi2 = 1536.3
Equil:3000 of 5000, Atoms:63(71)  Chi2 = 1497.57
Equil:4000 of 5000, Atoms:77(68)  Chi2 = 1529.76
Equil:5000 of 5000, Atoms:68(74)  Chi2 = 1513.21
Samp: 1000 of 5000, Atoms:62(81)  Chi2 = 1511.96
Samp: 2000 of 5000, Atoms:54(80)  Chi2 = 1492.69
Samp: 3000 of 5000, Atoms:59(83)  Chi2 = 1497.35
Samp: 4000 of 5000, Atoms:56(81)  Chi2 = 1498.62
Samp: 5000 of 5000, Atoms:65(83)  Chi2 = 1476.77
*** Check value of final chi2: 1476.77 ****
```

```
> plotGAPS(results$Amean, results$Pmean, 'GSFigs')
```

```
null device
1
```

Figure ?? shows the results from running CoGAPS on the GIST data in (?) with the option `plot` set to true.

Moreover, the gene set activity is provided in `results$GSActEst` including p-values for upregulation in `results$GSUpreg` and downregulation in `results$GSDownreg`.

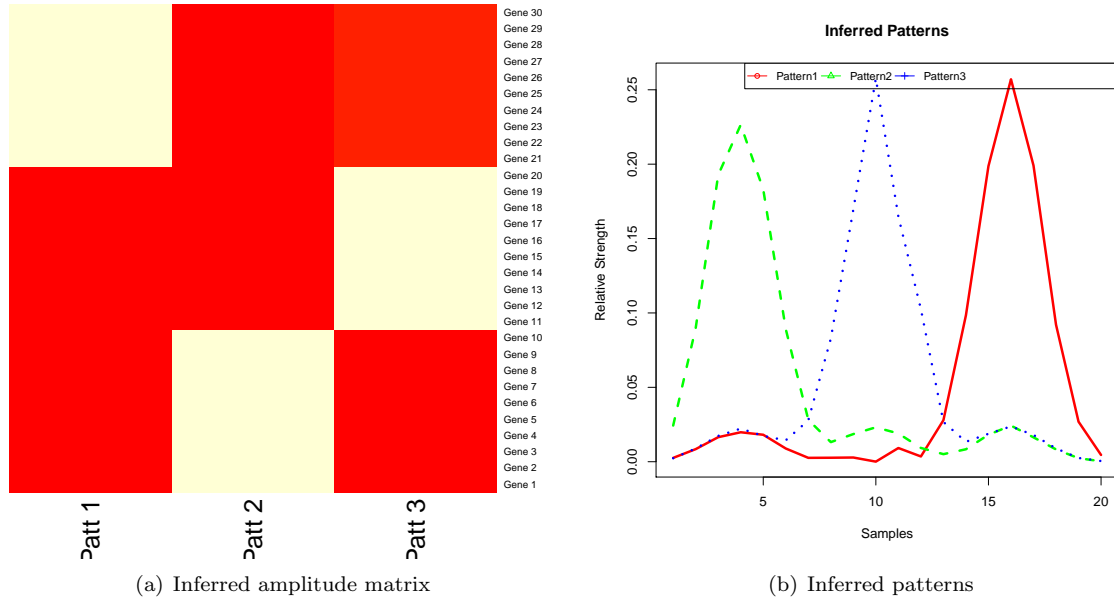


Figure 3.3: Results from GAPS on data of simulated gene set data.

3.2.3 Example: GIST Analysis

We also provide the code that repeats the CoGAPS analysis of GIST data (GIST_TS_20084) with gene sets defined by transcription factors (TFGSList), as in the DESIDE analysis of (?). To minimize time in running the vignette, this is not done live. However, running the code will complete the analysis and require on the order of an hour on a fast machine.

```
> library('CoGAPS')
> data('GIST_TS_20084')
> data('TFGSList')
> nIter <- 10000
> results <- CoGAPS(GIST.D, GIST.S, tf2ugFC,
+                   nFactor=5,
+                   nEquil=nIter, nSample=nIter,
+                   plot=FALSE)
> plotGAPS(results$Amean, results$Pmean, 'GISTFigs')
```

Figure ?? shows the results from running CoGAPS on the GIST data in (?) with the option plot set to true.

Moreover, the gene set activity is provided in `results$GSActEst` including p-values for upregulation in `results$GSUpreg` and downregulation in `results$GSDownreg`.

3.3 Using CoGAPS-based statistics to infer gene membership in annotated gene sets

As we describe in the previous section, the GAPS matrix factorization can be used to infer gene set activity in each pattern using the function `calcCoGAPSStat` (?). The `computeGeneGSProb` function extends this gene set statistic to compute a statistic quantifying the likely membership of each gene annotated to a set based




Figure 3.4: Results from GAPS on data of simulated gene set data.

upon its inferred activity (?). This statistic is formulated by comparing the expression pattern computed with CoGAPS of a given gene g annotated as a member of \mathcal{G} to the common expression pattern of all annotated members of \mathcal{G} . This similarity is quantified with the following summary statistic

$$S_{g,\mathcal{G}} = \frac{\sum_p -\log(\text{Pr}_{\mathcal{G}_p}) \mathbf{A}_{gp} / \mathbf{Asd}_{gp}}{\sum_p -\log(\text{Pr}_{\mathcal{G}_p})}, \quad (3.3)$$

where $\text{Pr}_{\mathcal{G}_p}$ is the probability of upregulation of the geneset, returned from `calcCoGAPSStat` as `GSAcEst` based upon eq. (??). Similar to the gene set statistics, p-values for the set membership are computed with permutation tests that compare the value of $S_{g,\mathcal{G}}$ from eq. (??) to the statistic formulated for a random gene set of the same size that also contains gene g .

The set membership statistic is computed from the results from the GAPS matrix factorization, computed with either the `GAPS` function described in Section ?? or the `CoGAPS` function described in Section ?? as follows:

```
> computeGeneGSProb(Amean, Asd, GStoGenes, numPerm=500)
```

Input Arguments

Amean Mean of the amplitude matrix estimated from the GAPS matrix factorization

Asd Standard deviation of the amplitude matrix estimated from the GAPS matrix factorization

GStoGenes List or data frame containing the genes in each gene set. If a list, gene set names are the list names and corresponding elements are the names of genes contained in each set. If a data frame, gene set names are in the first column and corresponding gene names are listed in rows beneath each gene set name.

nPerm Number of permutations used for the null distribution in the gene set and set membership statistics. (optional; default=500).

Function Output

p-value of set membership for each gene specified in `GStoGenes`.

3.3.1 Example on GIST Data

Although not run in the interest of installation time, the following examples were used to generate some of the results in (?), with the complete analysis code available from <http://astor.som.jhmi.edu/~ejfertif/Publications.html>.

This example refines transcription factor targets annotated in TRANSFAC (TFGSList) to identify context-specific targets from gene expression data (GIST-TS_20084) from (?).

```
> # define transcription factors of interest based on Ochs et al. (2009)
> TFs <- c("c.Jun", 'NF.kappaB', 'Smad4', "STAT3", "Elk.1", "c.Myc", "E2F.1",
+         "AP.1", "CREB", "FOXO", "p53", "Sp1")
> # take the results from the previously run analysis
> # set membership statistics
> permTFStats <- list()
> for (tf in TFs) {
+   genes <- levels(tf2ugFC[,tf])
+   genes <- genes[2:length(genes)]
+   permTFStats[[tf]] <- computeGeneTFProb(Amean = results$Amean,
+                                           Asd = results$Asd, genes)
+ }
```

Chapter 4

Feedback

Please send feedback to Elana Fertig ejfertig@jhmi.edu or Michael Ochs ochsm@tcnj.edu.

If you want to send a bug report, please first try to reproduce the error. The code is stochastic and we have seen many transient errors arising in the boost libraries which rarely repeat. Send the data and please describe what you think should have happened, and what did happen.

Chapter 5

Acknowledgments

This new version of CoGAPS relies on extensive recoding by Waisheng Lee in C++ and work for C++ streamlining and R integration by Ondrej Maxian and Conor Kelton. John Stansfield provided additional R code for visualization. Special thanks to paper co-authors Jie Ding, Alexander V. Favorov, and Giovanni Parmigiani for statistical advice in developing the GAPS / CoGAPS algorithms. Additional thanks to Simina M. Boca, Ludmila V. Danilova, Genevieve Stein-O'Brien, Jeffrey Leek, and Svitlana Tyekcheva for their useful feedback.

This work was funded by NIH/NLM R21LM009382, NIH/NLM R01LM011000, NIH/NCI K25CA141053, and NSF Grant 0342111.